Social Learning Equilibria

Elchanan Mossel[*], Manuel Mueller-Frank[†], Allan Sly[‡] and Omer Tamuz[§]

February 15, 2018

Abstract

We consider social learning settings in which a group of agents face uncertainty regarding a state of the world, observe private signals, share the same utility function, and act in a general dynamic setting. We introduce Social Learning Equilibria, a static equilibrium concept that abstracts away from the details of the given dynamics, but nevertheless captures the corresponding asymptotic equilibrium behavior. We establish strong equilibrium properties on agreement, herding, and information aggregation.

Keywords: Social learning; Agreement; Herding.
JEL: D83, D85.

1 Introduction

Social learning refers to the inference individuals draw from observing the behavior of others to their underlying private information. This inference then in turn impacts their own behavior. Social learning has served as an explanation for economic phenomena such as herding[1], bubbles and crashes in financial markets[2], optimal contracting[3], technology adoption[4] and more.[5]

[*]MIT.
[†]IESE.
[‡]Princeton University.
[§]Caltech. This work was supported by a grant from the Simons Foundation (#419427, Omer Tamuz).

[1]See Banerjee (1992); Bikhchandani et al. (1992).

[2]E.g., Scharfstein and Stein (1990); Welch (1992); Chari and Kehoe (2003).

[3]E.g., Khanna (1998); Arya et al. (2006).

[4]E.g., Walden and Browne (2002); Duan et al. (2009).

[5]Further references can be found in Bikhchandani et al. (1998), Chamley (2004), Vives (2010) and Jackson (2011).

Most theoretical contributions to social learning by rational agents have so far been based on a given dynamic game, which first specifies the *social learning setting*, i.e., the players, their actions and common utility function, the state and signal spaces, and a commonly known probability distribution thereover, and additionally specifies the extensive form of the game, including the order and frequency of decisions among players, and what each player knows at every given decision instant. This approach has two inherent weaknesses. First, the analysis of asymptotic equilibrium behavior in dynamic games is not straightforward, resulting in a limited range of tractable models and a focus on extremely stylized settings. Second, when trying to understand or predict behavior in real world social learning settings, the modeler might not know the exact nature of interaction among individuals, and the importance of each of the modeling assumptions is often unclear.

An important example is the sequential learning model of Bikhchandani et al. (1992), in which agents have to decide whether to (say) adopt or not adopt a new product. They receive conditionally i.i.d. private signals, and each makes a decision in an exogenously determined order, after observing the choices of their predecessors. In this highly stylized setting many interesting results were proved regarding herding and learning, with a particularly important contribution by Smith and Sørensen (2000). However, it is natural to wonder what happens in more realistic settings. What if the agents come in groups that act together? Perhaps they exchange information with the people standing behind them or in front of them in line? Perhaps they are allowed to change their decision in the five periods following the first one in which they acted?

We propose a different approach that abstracts away from the interaction structure, and focuses directly on the asymptotic steady state to which the equilibrium dynamic converges. We assume that each agent knows his private signal and additionally has some information about the signals of other agents, which she has presumably learned through some dynamic interaction. In a *social learning equilibrium* (SLE), which is a static concept, each agent chooses an action that is expected utility maximizing conditional on her information. One particular social learning equilibrium we analyze is the *complete social learning equilibrium* (CSLE) where each agent knows her private signal and the strategies and chosen actions of all other agents. This is a particularly natural equilibrium notion as it simply combines the concept of Nash equilibrium with that of social learning. Imposing mild conditions assures the existence of an SLE, and in fact of a CSLE.

The value of this approach stems from the connection between SLEs and the equilibria of generally defined social learning games. Theorem 6, which is essentially a reformulation of a result due to Rosenberg et al. (2009), establishes that asymptotic behavior in any Nash equilibrium[6] of any social learning game is captured by an SLE. This result implies that to understand the asymptotic equilibrium behavior of any social learning game it is sufficient to analyze the set of corresponding SLEs, greatly reducing the burden of equilibrium analysis.

For our further results we focus attention on the *canonical model* of social learning with countably many agents (i.e., a large group of agents) binary states, a common prior, conditionally i.i.d. signals[7], binary actions and a common utility function where the utility of each agent depends only on his action and the state of the world. For canonical settings we show that a number of phenomena (i.e., agreement, herding and information aggregation) that have been shown to emerge in the asymptotic states of particular dynamic social learning games, in fact appear very generally, as a property of the SLE. We further discover new connections between these phenomena.

Theorem 1 shows that every CSLE in a canonical setting satisfies *agreement*, i.e., all agents select the same action. The probability of the agreement action being correct can be linked to the structure of private signals. As defined by Smith and Sørensen (2000), private signals are unbounded if the support of the probability of either state conditional on one signal contains both zero and one. Theorem 2 shows that every CSLE in a canonical setting with unbounded signals aggregates all private information, i.e., the agreement action is optimal conditional on the realized state.

For the analysis of bounded signals, where the support of the belief conditional on one signal contains neither zero nor one, we borrow the concept of *information diffusion* introduced by Lobel and Sadler (2015) in the context of the sequential social learning model. Assume for simplicity that the support of private beliefs is $[\beta, 1-\beta]$. The action of an agent satisfies information diffusion if it is optimal given the state with a probability of at least $1-\beta$. Theorem 3 establishes that every CSLE in the canonical setting satisfies information diffusion. As for unbounded signals information diffusion implies information aggregation, Theorem 2 is thus a corollary of Theorem 3.

We next provide a general sufficient condition for herding, one of the most prominent concepts in the social learning literature. We do so in a more general setting in which we relax the assumption of conditional i.i.d. signals. Instead we assume that signals satisfy a *mixing* property. Informally, this means that conditional on the state, each agent's signal is almost independent of almost all the other agents' signals. A canonical* setting is a canonical setting with mixing signals. We say that *herding* occurs in an SLE if almost

[6]The set of Nash equilibria include the perfect Bayesian equilibria—whatever their definition might be in this case.

[7]For some proof we will require the additional, technical condition that the Kullback-Leibler divergence of the conditional measures is finite.

surely all agents but a finite subset select the same action. We say that an SLE is *weakly ordered* if there exists a weak order on the set of agents such that if $i \leq j$ then j knows i's action. Theorem 5 shows that every weakly ordered SLE satisfies herding.

Finally, Theorem 4 shows that in the canonical* setting in every SLE that satisfies herding the herding action satisfies information diffusion. This highlights a deep connection between the phenomena of agreement and of learning: when agents exchange enough information to agree on actions, they must in fact exchange a very large amount of information. Indeed, when signals are unbounded, they must exchange enough information to learn the state.

The social learning literature is too large to comprehensively cite here. We limit the discussion to those papers whose results are most closely related. Our equilibrium approach is more in line with Aumann's approach (1976) of studying a static environment with common knowledge, as compared to later social learning papers (e.g., Geanakoplos and Polemarchakis, 1982), which analyze the process by which common knowledge is reached. Similarly to Aumann, we directly study the equilibrium, rather than specifying the exact interaction structure and procedure by which the equilibrium is obtained. Our equilibrium notion is conceptually very closely related to that of a rational expectations equilibrium (Grossman, 1981). In its original formulation (for example Radner, 1979) the concept of rational expectation equilibrium (henceforth REE) is applied to market environments where participants have private information. A forecast function maps signal vectors into a pricing vector which is commonly observed by all agents. A REE is then a forecast function such that markets clear and for (almost) all signal realizations the portfolios of agents maximize their expected utility conditional on the forecast function and their private signal.

The main difference between our notion of social learning equilibrium and REE is threefold. First, we differ from their particular form of forecast function, which imposes a single summary statistic that is commonly observed. Second, under REE not only do actions have to be individually optimal as in our setting, but they additionally have to satisfy a market clearing condition. This difference arises from the fact that in our social learning setting payoff externalities are absent, contrary to a market environment. Third and most importantly, we show how this static equilibrium notion can serve to understand asyptotic equilibrium behavior in dynamic social learning games.

Minehart and Scotchmer (1999) introduce a concept of REE in a particular social learning setting. Despite some superficial similarities, their approach is essentially different from ours. For example, an equilibrium—as they define it—does not usually exist, and so they revert to an approximate equilibrium notion, in which they prove their main results.

The seminal paper by Bikhchandani et al. (1992) introduced the sequential social learning model where, in the canonical setting, agents make a one time choice observing the action chosen by all predecessors. They showed that a herd on the suboptimal action might emerge. Our Theorem 5 shows that herding is a much general feature of interaction among rational agents. Smith and Sørensen (2000) showed that the herding action in the sequential social learning model is almost surely optimal if signals are unbounded, and suboptimal with positive probability if signals are bounded. Our Theorem 4 shows that any herd by rational agents satisfies information diffusion. Thus the relation between herding and information aggregation established by Smith and Sørensen (2000) is an extremely robust outcome of interaction among rational individuals.

Our Theorem 3 is closely related to Lobel and Sadler (2015) who consider the sequential social learning model where each agent observes a random (possibly correlated) subset of her predecessors. They introduce the notion of information diffusion that we borrow for our analysis and provide two sufficient conditions on the random observation structure such that information diffuses respectively fails to diffuse in any equilibrium. Applying our results to their model we shed additional insight by connecting information diffusion to herding. Theorem 4 together with Theorem 6 implies that information diffuses in any random observation structure where herds occur with probability one.

Finally, our Theorem 1 extends the agreement results for settings of repeated interaction of Gale and Kariv (2003), Mueller-Frank (2013), and of Rosenberg et al. (2009) which all show that agreement occurs but in case of indifference among actions. Our Theorem 1 shows that in the canonical settings in a CSLE indifference occurs with probability zero.

The rest of the paper is organized as follows. Section §2 introduces the model and our equilibrium notion. Section §3 presents our results on agreement and information diffusion in CSLEs. Section §4 establishes our results on herding and information diffusion in SLEs. Section §5 establishes the formal relation between social learning equilibria and asymptotic equilibrium behavior in social learning games. Section §6 concludes.

2 The Model

We consider a group of agents who must each choose an action under uncertainty about a state of nature. Each agent's utility depends only on her own action and the state, and agents are homogeneous in the sense of sharing the same utility function. Each agent observes a private signal, and additionally some information about the others' signals. A social learning equilibrium is simply a choice of action for each agent that maximizes her expected utility, given the information available to her; note that this information may include the choices of others. We now define this formally.

Social learning settings

A *social learning setting* $(N, A, \Theta, u, S, \mu)$ is defined by a set of players N, a compact metrizable (common) action space A, a compact metrizable state space Θ, a continuous utility function $u : A \times \Theta \to \mathbb{R}$, a measurable private signal space S, and finally a commonly known joint probability distribution μ over $\Theta \times S^N$.

We will denote by θ the random state of nature and by $\bar{s} = (s_i)_{i \in N}$ the agents' private signals. When no ambiguity arises we will denote probabilities and expectations with respect to μ by $\mathbb{P}[\cdot]$ and $\mathbb{E}[\cdot]$, respectively. For some modeling applications it will furthermore be useful to add to this probability space a non-atomic random variable r that is independent of the rest.

Social learning equilibria (SLE)

Each agent i, in addition to her private signal s_i, learns ℓ_i, which is some function of $\bar{s}$ (and possibly r). Agent i's (random) action is a_i. It takes values in A, and is some function of ℓ_i and s_i. Equivalently, ℓ_i and a_i are random variables that are, respectively, $\sigma(\bar{s}, r)$- and $\sigma(\ell_i, s_i)$-measurable.

Let $\bar{\ell}$ and $\bar{a}$ denote $(\ell_i)_{i \in N}$ and $(a_i)_{i \in N}$, respectively. In a given social learning setting, a *social learning equilibrium* (or SLE) is a pair $(\bar{\ell}, \bar{a})$ such that each agent's action a_i is a best response, given her information ℓ_i and s_i:

$$a_i \in \operatorname*{argmax}_{a \in A} \mathbb{E}[u(a, \theta) \mid \ell_i, s_i]. \tag{1}$$

Complete social learning equilibria (CSLE)

The first class of social learning equilibria which we consider are *complete social learning equilibria* (or CSLE). In a CSLE $\ell_i = \bar{a}$. That is, each agent, in addition to her private signal, learns the actions of all other agents. Thus, in a given social learning setting, the actions $\bar{a}$ are a CSLE if it holds that

$$a_i \in \operatorname*{argmax}_{a \in A} \mathbb{E}[u(a, \theta) \mid \bar{a}, s_i]. \tag{2}$$

To specify a CSLE it suffices to specify the actions $\bar{a}$, since $\ell_i = \bar{a}$ for all i.

Note that a related, natural and more general class of SLEs are those in which $\bar{a}$ is $\sigma(\ell_i, s_i)$-measurable. That is, those SLEs in which the agents all know each other's actions, and perhaps more information additionally. We will prove our results on CSLEs in this generality, but prefer to adhere to the definition above because of its simplicity and proximity to Nash equilibria.

Existence

In every social learning setting there exists an SLE, and moreover a CSLE. This follows directly from the existence of an optimal action, given knowledge of all the private signals. For a CSLE that always exists, let $a^* = a^*(\bar{s})$ be an action that maximizes expected utility conditional on $\bar{s}$, the entire collection of private signals, and set $a_i = a^*$ for all $i \in N$. As a^* aggregates all private information we call such an equilibrium *information aggregating*.

3 Agreement and Information Aggregation in Complete Social Learning Equilibria

In this section we study complete social learning equilibria (CSLEs). We focus on a class of social learning settings which appears frequently in the literature: in *canonical settings* N is countably infinite, $A = \Theta = \{0,1\}$, $u(1,1) = u(0,0) = 1$ and $u(0,1) = u(1,0) = 0$, and signals are informative and conditionally i.i.d.

Agreement

An SLE satisfies *agreement* if almost surely we have $a_i = a_j$ for all pairs of agents i, j. Our first result establishes agreement as a property of any CSLE.

Theorem 1. *In a canonical setting every CSLE satisfies agreement.*

This result shows that Aumann's seminal agreement result carries over to canonical social learning settings as a property of every complete social learning equilibrium. The conceptual reasoning behind the result, however, differs as no epistemic conditions beyond knowledge of the information structure and the social learning equilibrium are required. Previous results in the literature have established that agreement is achieved, except in cases of indifference (Mueller-Frank, 2013; Rosenberg et al., 2009). Our contribution is to show that, for the case of CSLE in canonical settings, indifference almost surely does not occur and hence agreement holds almost surely. For finite but large groups we show in the appendix that agreement holds with high probability in any CSLE.

To prove this result we first show that whenever both actions are taken, it must be that all agents are indifferent between the actions. This follows from the same intuition that underlies the no trade theorem of Milgrom and Stokey (1982), as well as and similar results in social learning (e.g., Sebenius and Geanakoplos, 1983; Mueller-Frank, 2013; Rosenberg et al., 2009). Next, we show that when all agents are indifferent then their belief must equal the belief that they would have if they revealed all of their private signals. The proof of this part follows the proof of a similar result for rational expectations equilibria by DeMarzo and Skiadas (1999b). Finally, since revealing all the private information cannot result in exact indifference—for an infinite group of agents—we conclude that this happens with probability zero. Likewise, this happens with probability zero for finite groups with non-atomic signals, or with generic priors. For large finite groups this can only happen with a probability that vanishes to zero exponentially in the size of the group, for a fixed distribution of private signals (Corollary 2).

Information aggregation

We next turn to the question of the learning properties of the agreement equilibrium. Under which conditions is the agreement action optimal? In terms of our definitions from the previous section, we ask: under which conditions is every CSLE information aggregating?

The *private belief* p_i of an agent is equal to her posterior probability conditional on her private signal only:

$$p_i = \mathbb{P}[\theta = 1 \mid s_i]. \tag{3}$$

As defined by Smith and Sørensen (2000), private signals are *unbounded* if the support of the private belief contains both 0 and 1. Similarly, private signals are *bounded* if the support of private beliefs contains neither 0 nor 1. Smith and Sørensen (2000) showed that in the sequential social learning model unbounded private signals are sufficient for agents eventually to select the action that corresponds to the true state. The following result relates the unbounded signal property to information aggregation in social learning equilibria.

Theorem 2. *In a canonical setting with unbounded signals every CSLE is information aggregating.*

What can be said about the learning properties of the agreement action if signals are not unbounded? Since independent of the signal structure there always exists an information aggregating equilibrium, the question is what is the worst possible equilibrium outcome in terms of information aggregation. To answer this, we borrow the notion of information diffusion introduced by Lobel and Sadler (2015) in context of the sequential social learning model. Consider the support of the private belief and let its convex hull be $[\beta_L, \beta_H]$. For simplicity assume that the support is symmetric, i.e., $[\beta, 1-\beta]$. A CSLE satisfies *information diffusion* if the probability of any agent's action being optimal is at least $1-\beta$:

$$\mathbb{P}[a_i \in \operatorname*{argmax}_{a\in A} u(a,\theta)] \geq 1-\beta.$$

As Lobel and Sadler highlight, the notion of information diffusion is particularly insightful if strong signals, i.e., those that induce a posterior belief close to β or $1-\beta$, are rare. The following is our main result of this section.

Theorem 3. *In a canonical setting every CSLE satisfies information diffusion.*

Note that Theorem 3 directly implies Theorem 2, as for unbounded signals we have $\beta = 0$. Theorem 3 is in turn a corollary of Theorem 4; we discuss in §4 the intuition behind this result.

4 Herding, agreement and Information Diffusion

Arguably the most prominent result in the social learning literature is the herding result established by Bikhchandani et al. (1992) in the canonical sequential social learning model. They show that if agents make an irreversible binary decision in strict sequential order, observing all the actions taken before them, then eventually all agents take the same action, i.e., a herd occurs. Furthermore, the action chosen by this herd is not necessarily optimal, even though the information contained in the pooled private signals suffices to choose the optimal action. Later, Smith and Sørensen (2000) showed that when signals are unbounded a herd still occurs, but the action chosen by the herd is almost surely optimal.

In analogy to this herding phenomenon, we say that an SLE satisfies *herding* if there is almost surely a cofinite set of agents who choose the same action. We call this (random) action the *herding action*. We say that the herding action satisfies information diffusion if it is equal to the optimal action (given knowledge of the state) $\operatorname{argmax}_{a\in A} u(a,\theta)$ with probability at least $1-\beta$ when the convex hull of the support of private beliefs is $[\beta, 1-\beta]$. For unbounded signals this is equivalent to the herding action being equal to the optimal action, given the state.

We would like to emphasize that despite the image that the term "herd" evokes, herding does not imply that the agents take a mindless, suboptimal action; indeed, the action chosen by the herd can be correct with probability one, as in the case unbounded signals in the model of Smith and Sørensen (2000). Accordingly, we think of a herding as a form of weak *agreement*: a herding SLE in one in which almost all the agents agree.

We must learn to agree

Our first result of this section highlights a deep connection between agreement and learning: one cannot agree without learning. In other words, in order for agents to agree they must exchange a large amount of information, and in particular such a large amount that they learn much about the state in the process. We thus offer the phrase "we must learn to agree" not as an imperative but as an observation.

We prove this result in a slightly more general setting than the canonical setting. In particular, we relax the requirement that signals are i.i.d., and require only a form of *mixing*. This holds when (1) the distributions of each of the private signals s_i, conditional on the state, are all identical, and (2) conditionally on the state $\theta \in \{0,1\}$, for each measurable subset $S' \subseteq S$ of the private signal space and for each $\varepsilon > 0$, there is an N such that for each agent i there are at most N agents j such that

$$\Big|\mathbb{P}[s_i \in S', s_j \in S' \mid \theta] - \mathbb{P}[s_i \in S', \mid \theta] \cdot \mathbb{P}[s_j \in S' \mid \theta]\Big| > \varepsilon.$$

Intuitively, private signals are mixing when for each agent i there are only finitely many other agents with whom i has a significantly correlated signal. One obvious example of mixing signals are i.i.d. signals. Other examples arise naturally when agents who are close to each other—either geographically or temporally—observe same or similar signals, but agents who are far away from each other observe only very weakly related signals.

In *canonical* settings* N is countably infinite, $A = \Theta = \{0,1\}$, the utility function assigns 1 if the action of an agent matches the state and 0 otherwise, and signals are informative and mixing. The next theorem shows that in canonical* settings, "herding implies learning".

Theorem 4. *In a canonical* setting, and when signals are unbounded, then in every SLE that satisfies herding, the herding action satisfies information diffusion.*

Thus, in any social learning environment where herding occurs, it is necessary that a large amount of information is exchanged among agents. To see this, consider a setting with unbounded signals. Here for the herding action to be optimal it needs to embed the private signals of infinitely many agents. For the case of bounded signals, consider a private belief distribution where the probability mass is concentrated on a close neighborhood of $\frac{1}{2}$ and only very little mass on neighborhoods of the endpoints $\beta, 1-\beta$. Also in this case it is clear that for the herding action to satisfy information diffusion it needs to incorporate many of the private signals. Since agreement is a strong form of herding we can conclude that "agreement implies learning", and more generally that "herding implies learning.

The proof of this theorem relies on the following lemma, which distills the important feature of mixing signals. Informally, it states that every event that is measurable in the private signals is conditionally approximately independent of almost all of them.

Lemma 1. *Consider a canonical* setting. Let B be any event that is $\sigma(\bar{s}, r)$-measurable. Let S' be any subset of the private signal space. Then for every $\varepsilon > 0$ there are at most finitely many agents i such that*

$$\Big|\mathbb{P}[s_i \in S', B \mid \theta] - \mathbb{P}[s_i \in S' \mid \theta] \cdot \mathbb{P}[B \mid \theta]\Big| > \varepsilon.$$

Given this lemma, the proof of Theorem 4 has a simple intuition: consider the herding action of the SLE. By Lemma 1 it is approximately independent of almost all the private signals. Yet, it is taken by almost all the agents. Therefore, there will be an agent that takes the herding action with very high probability, and whose signal is almost completely independent from it. Hence such an agent would prefer to follow her own private signal whenever doing so is more likely to be correct than following the herd. But in equilibrium this agent does follow the herd, and so it must be that her private signals never give an indication that is stronger than the signal contained in the herding action.

This result is related to similar results for rational expectations equilibria (e.g., DeMarzo and Skiadas, 1999a). There, however, agreement implies efficient aggregation of information even for a small number of players and bounded signals, whereas in our setting this holds less generally and crucially depends on both the large size of the group and the unboundedness of the signals.

Herding in weakly ordered SLEs

We study a class of SLEs that correspond to a large class of social learning games, including the sequential models of Bikhchandani et al. (1992) and Smith and Sørensen (2000). We focus on canonical settings, and on SLEs $(\bar{\ell}, \bar{a})$ that are *weakly ordered*: there is some weak order $\leq$ on the agents such that, if $i \leq j$ then agent j observes i's action: a_i is $\sigma(\ell_j)$-measurable. The case that the order is strict and $\ell_j = (a_1, \ldots, a_{j-1})$ corresponds to the classical sequential models. Weakly ordered SLEs correspond to a much wider class of social learning games: perhaps the agents come in groups that act together; perhaps they exchange information by cheap talk with the people standing behind them or in front of them in line. Perhaps they are allowed to change their decision in the five periods following the first one in which they acted, etc.

We show that every weakly ordered SLE satisfies *herding*: there is almost surely a cofinite set of agents who choose the same action. We call this (random) action the *herding action*.

Theorem 5. *In a canonical* setting every weakly ordered SLE satisfies herding.*

A direct corollary of Theorems 4 and 5 is the following theorem, which is a generalization of the results of Smith and Sørensen (2000).

Corollary 1. *In a canonical* setting, and when signals are unbounded, then in every weakly ordered SLE the herding action satisfies information diffusion.*

The proof of Theorem 5 uses similar ideas to that of Theorem 4, but is somewhat more complicated. Here, one must first observe that if both actions are taken infinitely often then agents must asymptotically be indifferent. If this occurs with positive probability, then eventually agents will be able to guess (correctly with high probability) that this will happen. Since—again asymptotically—almost all agents have signals that are independent of this event, they would choose to ignore it and follow their own private signals. But then they would not be indifferent, and thus this cannot happen with positive probability.

5 Social Learning Equilibria and Social Learning Games

In this section we consider a large class of *social learning games*. Given a social learning setting, a social learning game is a dynamic game with incomplete information in which agents choose actions and observe information about other agents' actions and signals. This class includes many models studied in the literature, including sequential learning models and models of repeated interaction on social networks.

The main result of this section relates social learning games to SLEs. We show that the asymptotic equilibrium behavior of agents in any social learning game is captured by an SLE: for any distribution over asymptotic equilibrium action profiles of a social learning game there exists an SLE with a matching distribution over action profiles.

This correspondence provides motivation for studying SLEs, and also allows to understand the long-run behavior of agents in many dynamic settings.

Social learning games

A *social learning game* includes a social learning environment $(N, A, \Theta, u, S, \mu)$, together with a description of the dynamics by which agents interact and learn. This consists of the tuple $(T, \Bbbk, \sigma, \delta)$. For each agent i the set $T_i \subseteq \{1, 2, \ldots\}$ denotes the set of *action times* i, i.e., the set of time periods in which agent i exogenously "wakes up", receives information and takes an action. The set $T = (T_i)_{i \in N}$ denotes the tuple of action times. For each agent i and time $t \in T_i$, let $\Bbbk_{i,t}$ be the information learned by agent i at time t, and let $\sigma_{i,t}$ be the action taken by agent i at time t. We denote by

$$\Bbbk_i^t = \{\Bbbk_{i,\tau} : \tau \leq t, \tau \in T_i\}$$

the information observed by agent i by time t, and by

$$\Bbbk_i = \{\Bbbk_{i,t} : t \in T_i\}$$

all the information observed by her, excluding her signal. We denote by

$$\sigma_i^t = \{\sigma_{i,\tau} : \tau < t, \tau \in T_i\}$$

the actions taken by agent i *before* time t, and by

$$\sigma^t = (\sigma_i^t)_{i \in N}$$

all the actions taken by all the agents before time t.

Each action $\sigma_{i,t}$ takes values in A, and is some function of the information known to agent i at time t, which consists of $\Bbbk_i^t$ and her private signal s_i:

$$\sigma_{i,t} = \sigma_{i,t}(\Bbbk_i^t, s_i).$$

The collection of maps $\sigma_i = (\sigma_{i,t})_{t \in T_i}$ is player i's strategy, and the tuple of strategies across all agents, $(\sigma_i)_{i \in N}$, is the strategy profile.

The information $\Bbbk_{i,t}$ is some function of the agents' actions before time t, the private signals themselves, as well as the additional independent random variable r, and takes values in some measurable space:

$$\Bbbk_{i,t} = \Bbbk_{i,t}(\sigma^t, \bar{s}, r).$$

The (possible) dependence on r allows this framework to include mixed strategies.

Finally, δ is the common discount factor, and agent i's discounted expected utility is

$$\sum_{t \in T_i} \delta^t \cdot \mathbb{E}[u(\sigma_{i,t}, \theta)].$$

A strategy profile σ is a *Nash equilibrium* if for each agent i her strategy σ_i maximizes her discounted expected utility given σ_{-i}, among all possible strategies for player i.

If agents are myopic, i.e. $\delta = 0$, a strategy profile σ is a Nash equilibrium if for each agent i, given σ_{-i}, in each period t her strategy $\sigma_{i,t}$ maximizes her expected utility in period $t \in T_i$ conditional on $\Bbbk_i^t$ and s_i:

$$\sigma_{i,t} \in \operatorname*{argmax}_{a \in A} \mathbb{E}[u(\sigma_{i,t}, \theta) \mid \Bbbk_i^t, s_i].$$

This definition of a social learning game is rather general and captures a variety of different models. Most prominently it captures the sequential social learning model of Bikhchandani et al. (1992), and Smith and Sørensen (2000). To see this simply set $T_i = \{i\}$ for every agent i and

$$\Bbbk_{i,i} = \Bbbk_i = \{\sigma_{j,j} : j < i\}.$$

The sequential social learning models of Acemoglu et al. (2011), Lobel and Sadler (2015) and others are likewise included in this framework. Here we have $T_i = \{i\}$ again, but $\Bbbk_{i,t}$ does not include all the actions of the predecessors, but rather only those of a random subset of the predecessors. The models of repeated interaction on social networks of Gale and Kariv (2003), Mossel et al. (2014) and Mossel et al. (2015) can be captured by setting $T_i = \mathbb{N}$ for all agents i and where $\Bbbk_{i,t}$ contains the last period actions of all the neighbors of agent i. Rosenberg et al. (2009) study a more general model that is not subsumed by the framework, but still shares many similarities. In fact, the proof of our result for this section, Theorem 6, follows exactly the proof of their Proposition 2.1.

Finally, the models of repeated communication of beliefs in a social network analyzed in Geanakoplos and Polemarchakis (1982), Parikh and Krasucki (1990), and Mueller-Frank (2013) can be captured by a squared loss utility function and a discount factor equal to zero, hence inducing myopic behavior.

Let $\bar{A}_i$ denote the (random) set of *accumulation points* of agent i's realized actions $(\sigma_{i,t})_{t \in T_i}$; if T_i is finite, then let $\bar{A}_i$ be the singleton that contains only the last period action of agent i. If T_i is infinite and A finite, then $\bar{A}_i$ consists of the actions chosen infinitely often.

Given these definitions, we are ready to establish the relation between Nash equilibria of the social learning game and SLEs. As we mention above, this theorem is essentially due to Rosenberg et al. (2009).

Theorem 6. *Consider a social learning environment and a social learning game* $(T, \Bbbk, \sigma, \delta)$*. Let* $\bar{a}$ *be any (random) action such that* $a_i \in \bar{A}_i$*, and let* $\ell_i = \Bbbk_i$*. Then* $(\bar{\ell}, \bar{a})$ *is an SLE.*

Thus, the asymptotic behavior in any Nash equilibrium[8] of a social learning game is captured by an SLE.

Learning and agreeing in social learning games

Recall that in Theorem 6 we showed that the asymptotic outcomes of social learning games correspond to social learning equilibria. Therefore, a straightforward application of Theorem 4 to social learning games implies that in every social learning game which is played in a canonical* setting with unbounded signals, it holds that if herding occurs (i.e., if a cofinite set of agents *converges* to the same action) then a cofinite set of agents *converges* to the correct action. Likewise, Corollary 1 implies that herding is indeed the outcome across a large spectrum of social learning games: it suffices that there is a weak order on the agents such that if $i \leq j$ then j observes which actions i chooses infinitely often. This generalizes the results of Smith and Sørensen (2000), highlighting the deeper forces that drive them.

6 Conclusion

We introduced the concept of social learning equilibrium as a useful tool to analyze social learning. The advantage over the conventional approach is that results or predictions can be made without knowing the exact dynamic of the interaction structure. We provide the agreement, herding and information aggregation results of social learning equilibria that unify and shed additional insight on the literature on social learning. In particular, we show that the relation between unbounded signals and the optimality of the herding action established by Smith and Sørensen (2000) holds much more generally. In fact, in any social learning environment with unbounded signals the action selected in a Bayesian herd is optimal.

The main value deriving from our analysis, however, is to show that the asymptotic equilibrium behavior of any social learning game might be analyzed via the static concept of social learning equilibrium, greatly simplifying the analysis.

[8] The set of Nash equilibria include the perfect Bayesian equilibria—whatever their definition might be in this case.

Finally, our concept can easily be adjusted to capture payoff heterogeneity and payoff externalities. This extension is kept for future work.

References

Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011.

Anil Arya, Jonathan Glover, and Brian Mittendorf. Hierarchical reporting, aggregation, and information cascades. *Managerial and Decision Economics*, 27(5):355–362, 2006.

R.J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976. ISSN 0090-5364.

Abhijit V Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, pages 797–817, 1992.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, pages 992–1026, 1992.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of economic perspectives*, 12(3):151–170, 1998.

Christophe Chamley. *Rational herds: Economic models of social learning*. Cambridge University Press, 2004.

Varadarajan V Chari and Patrick J Kehoe. Hot money. *Journal of Political Economy*, 111 (6):1262–1292, 2003.

Peter DeMarzo and Costis Skiadas. On the uniqueness of fully informative rational expectations equilibria. *Economic Theory*, 13(1):1–24, 1999a.

Peter DeMarzo and Costis Skiadas. On the uniqueness of fully informative rational expectations equilibria. *Economic Theory*, 13(1):1–24, 1999b.

Wenjing Duan, Bin Gu, and Andrew B Whinston. Informational cascades and software adoption on the internet: an empirical investigation. *MIS quarterly*, pages 23–48, 2009.

Douglas Gale and Shachar Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, 2003.

John D Geanakoplos and Heraklis M Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982.

Sanford J Grossman. An introduction to the theory of rational expectations under asymmetric information. *The Review of Economic Studies*, 48(4):541–559, 1981.

Matthew O Jackson. An overview of social networks and economic applications. In *Handbook of social economics*, volume 1, pages 511–585. Elsevier, 2011.

Naveen Khanna. Optimal contracting with moral hazard and cascading. *The Review of Financial Studies*, 11(3):559–596, 1998.

Ilan Lobel and Evan Sadler. Information diffusion in networks through social learning. *Theoretical Economics*, 10(3):807–851, 2015.

Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of economic theory*, 26(1):17–27, 1982.

D. Minehart and S. Scotchmer. Ex post regret and the decentralized sharing of information. *Games and Economic Behavior*, 27(1):114–131, 1999.

Elchanan Mossel, Allan Sly, and Omer Tamuz. Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields*, 158(1-2):127–157, 2014.

Elchanan Mossel, Allan Sly, and Omer Tamuz. Strategic learning and the topology of social networks. *Econometrica*, 83(5):1755–1794, 2015.

Manuel Mueller-Frank. A general framework for rational learning in social networks. *Theoretical Economics*, 8(1):1–40, 2013.

R. Parikh and P. Krasucki. Communication, consensus, and knowledge. *Journal of Economic Theory*, 52(1):178–189, 1990. ISSN 0022-0531.

Roy Radner. Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica: Journal of the Econometric Society*, pages 655–678, 1979.

D. Rosenberg, E. Solan, and N. Vieille. Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979–994, 2009. ISSN 0899-8256.

David S Scharfstein and Jeremy C Stein. Herd behavior and investment. *The American Economic Review*, pages 465–479, 1990.

James K Sebenius and John Geanakoplos. Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424–426, 1983.

Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.

Xavier Vives. *Information and learning in markets: the impact of market microstructure*. Princeton University Press, 2010.

Eric Walden and Glenn Browne. Information cascades in the adoption of new technology. *ICIS 2002 Proceedings*, page 40, 2002.

Ivo Welch. Sequential sales, learning, and cascades. *The Journal of finance*, 47(2):695–732, 1992.

A Agreement and Information Aggregation in Complete Social Learning Equilibria

To prove Theorem 1 we prove the following, stronger result, which applies to CSLEs in canonical settings, as well as CSLEs in canonical settings but with finitely many agents. This proof uses similar ideas to the proof of Proposition 3 of DeMarzo and Skiadas (1999b).

Theorem 7. *Let $(\bar{\ell}, \bar{a})$ be a SLE where $\bar{a}$ is $\sigma(\ell_i, s_i)$-measurable for every i, in a canonical setting, but with either finitely many or countably infinitely many agents.*

Let D be the event that there is disagreement, i.e.,

$$D = \{a_i \neq a_j \text{ for some } i, j\}.$$

Let

$$q_s = \mathbb{P}[\theta = 1 \mid \bar{s}]$$

be the belief induced by the collection of all private signals. Then the event $\{q_s = 1/2\}$ contains the event D (mod 0), and in particular

$$\mathbb{P}[D] \le \mathbb{P}[q_s = 1/2].$$

The following is an immediate corollary of Theorem 7, since $\mathbb{P}[q_s = 1/2] = 0$ whenever N is finite, generically, and when private beliefs are non-atomic, and since, by the Chernoff bound, this probability decays exponentially in $|N|$, for a fixed private signal distributions.

Corollary 2. *In the setting of Theorem 7, under any of the following conditions, every equilibrium satisfies agreement.*

1. *When N is infinite.*
2. *Generically, over the choice of prior and private signal distributions.*
3. *When private beliefs are non-atomic.*

Furthermore, for any fixed private signal distribution, $\mathbb{P}[D]$ decreases exponentially in the number of agents. Thus the agents all agree, except with probability that diminishes exponentially.

In particular the setting in which N is infinite is the setting of Theorem 1, which is therefore a corollary of Theorem 7.

To prove this theorem we will need the following standard lemma, which is a form of the "No Trade Theorem" of Milgrom and Stokey (1982). Denote agent i's equilibrium belief by

$$q_i = \mathbb{P}[\theta = 1 \mid \ell_i, s_i].$$

Lemma 2. *In the setting of Theorem 7, if D has positive probability, then conditioned on D it almost surely holds that $q_i = 1/2$ for all i.*

Proof. Consider an outside observer who observes all the agents' actions $\bar{a}$. Her belief is

$$q_* = \mathbb{P}[\theta = 1 \mid \bar{a}].$$

Since $\bar{a}$ is $\sigma(\ell_i, s_i)$-measurable, it follows from the law of total expectations that for every i

$$q_* = \mathbb{E}[q_i \mid \bar{a}]. \tag{4}$$

Since 1 is the action that is optimal for beliefs above 1/2, we have that $a_i = 1$ implies that $q_i \geq 1/2$. Likewise, $a_i = 0$ implies $q_i \leq 1/2$. Hence the claim follows by (4). ∎

Denote by μ_1 and μ_0 the distributions of each private signal s_i, conditioned on $\theta = 1$ and $\theta = 0$, respectively. Denote agent i's *private log-likelihood ratio* by

$$z_i = \log \frac{d\mu_1^i}{d\mu_0^i}(s_i),$$

and, as in (3), let

$$p_i = \mathbb{P}[\theta = 1 \mid s_i]$$

be the belief induced by agent i's private signal only.

If we denote

$$z_0 = \log \frac{\mathbb{P}[\theta = 1]}{\mathbb{P}[\theta = 0]}$$

then

$$\log \frac{p_i}{1 - p_i} = z_0 + z_i.$$

Denote $z = \lim_{n \to |N|} \frac{1}{n} \sum_{i=0}^{n} z_i$; here z takes values in $\mathbb{R}$. Note that even when $|N| = \infty$ then the limit a.s. exists by the strong law of large numbers, as the sequence $(z_i)_i$ is independent, conditioned on θ. Here we use the assumption that the Kullback-Leibler divergence of μ_0 and μ_1 is finite; this is equivalent to the finiteness of the expectation of z_i.

Since the private signals are conditionally independent, it follows by Bayes' rule that

$$q_s = L(z), \tag{5}$$

where

$$L(x) = \frac{e^{|N|x}}{1 + e^{|N|x}},$$

when N is finite, and

$$L(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$$

when N is infinite.

Fix any agent i, and let

$$q_i = \mathbb{P}[\theta = 1 \mid \bar{a}, s_i]$$

be the equilibrium belief of agent i. By Lemma 2, conditioned on D it holds almost surely that $q_i = 1/2$.

Now,

$$q_i = \mathbb{P}[\theta = 1 \mid q_i] = \mathbb{E}\left[\mathbb{P}[\theta = 1 \mid \bar{s}] \mid q_i\right] = \mathbb{E}\left[q_s \mid q_i\right], \tag{6}$$

where the first equality follows from the law of total expectations (as q_i is a function of the private signals), and the second from the definition of q_s.

Our goal is to show that conditioned on D, $q_s = 1/2$. Using Lemma 2, we can show this by showing that $q_s = q_i$, conditioned on D. To this end, we will show that conditioned on q_i and D, z and $L(z)$ are linearly independent. It will follow that conditioned on q_i and D, z is constant, and therefore, by (6), $q_s = q_i = L(z) = 1/2$.

By the law of total expectation we have that

$$\begin{aligned} \mathbb{E}[z_i \cdot q_s \mid D] &= \mathbb{E}[\mathbb{E}[z_i \cdot q_s \mid z_i, D] \mid D] \\ &= \mathbb{E}[z_i \cdot \mathbb{E}[q_s \mid z_i, D] \mid D]. \end{aligned}$$

Here we again use the assumption that the expectation of z_i is finite.

Since $q_i = 1/2$ conditioned on D, since z_i is a function of s_i, and by (6), $\mathbb{E}[q_s \mid z_i, D] = 1/2$. Hence

$$\mathbb{E}[z_i \cdot q_s \mid D] = \mathbb{E}[z_i \mid D] \cdot \frac{1}{2}.$$

Again using the facts that $q_i = 1/2$ conditioned on D and (6), we have that $\mathbb{E}[q_s \mid D] = 1/2$. Recalling that $q_s = L(z)$, we have shown that

$$\mathbb{E}[z_i \cdot L(z) \mid D] = \mathbb{E}[z_i \mid D] \cdot \mathbb{E}[L(z) \mid D]. \tag{7}$$

We now consider two cases. First, assume that N is finite. Then summing (7) over $i = 0, \ldots, |N|$ and dividing by $|N|$ yields

$$\mathbb{E}[z \cdot L(z) \mid D] = \mathbb{E}[z \mid D] \cdot \mathbb{E}[L(z) \mid D].$$

Now, since $L(z)$ is a monotone function of z, by Chebyshev's sum inequality we have that

$$\mathbb{E}[z \cdot L(z) \mid D] \geq \mathbb{E}[z \mid D] \cdot \mathbb{E}[L(z) \mid D].$$

with equality only if z (or, equivalently $L(z)$) is constant. Hence z is constant conditioned on D and the proof of Theorem 7 is concluded, for the case that $|N| < \infty$.

Consider now the case that N is infinite. We apply the same argument, but notice first that

$$\lim_{n\to\infty} \mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n} z_i \cdot L(z) \middle| D\right] = \mathbb{E}\left[\lim_{n\to\infty} \frac{1}{n}\sum_{i=0}^{n} z_i \cdot L(z) \middle| D\right] = \mathbb{E}[z \cdot L(z) \mid D],$$

since the sequence of partial averages of $(z_i)_i$ converges almost surely and in L^1 (by the strong law of large numbers), and since $L(z) \cdot \mathbb{1}_{\{D\}}$ is bounded in $[0,1]$. Likewise

$$\lim_{n\to\infty} \mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n} z_i \mid D\right] = \mathbb{E}[z \mid D],$$

and so we can indeed proceed with the argument as in the case of finite N. This completes the proof of Theorem 7.

B Herding

In this section we prove Theorem 5.

Let

$$x_i = \mathbb{P}[\theta = 1 \mid a_1, \ldots, a_i]$$

be the sequence of *public beliefs*, and let

$$q_i = \mathbb{P}[\theta = 1 \mid \ell_i, s_i]$$

be agent i's equilibrium belief. Note that, since each agent i knows $\{a_1, \dots, a_i\}$,

$$x_i = \mathbb{E}[q_i \mid a_1, \dots, a_i], \tag{8}$$

by the law of total expectations.

Note that the action 1 is optimal for beliefs 1/2 and higher, and the action 0 is optimal for beliefs 1/2 and lower. Therefore, and since $\bar{a}$ is an equilibrium,

$$a_i = 1 \Rightarrow q_i \geq 1/2 \quad \text{ and } a_i = 0 \Rightarrow q_i \leq 1/2 \tag{9}$$

and

$$\mathbb{E}[u(a_i, \theta) \mid q_i] = \mathbb{P}[a_i = \theta \mid q_i] = \max\{q_i, 1 - q_i\}. \tag{10}$$

We start with two simple claims regarding a_i and x_i.

Claim 1. *If $a_i = 1$ then $x_i \geq 1/2$. If $a_i = 0$ then $x_i \leq 1/2$.*

Proof. By (9) we have that $q_i \geq 1/2$ conditioned on $a_i = 1$. Hence, by (8), $x_i \geq 1/2$ conditioned on $a_i = 1$. An analogous argument holds for the case $a_i = 0$. ∎

Claim 2. $\mathbb{P}[a_i = \theta \mid x_i] = \max\{x_i, 1 - x_i\}$.

Proof. By Claim 1

$$\mathbb{P}[\theta = a_i \mid x_i] = \begin{cases} \mathbb{P}[\theta = 1 \mid x_i] & \text{if } x_i > 1/2 \\ \mathbb{P}[\theta = 0 \mid x_i] & \text{if } x_i < 1/2 \\ \mathbb{P}[\theta = a_i \mid x_i] & \text{if } x_i = 1/2. \end{cases}$$

By (8) and (9), if $x_i = 1/2$ then $x_i = q_i$. Therefore, and since $\mathbb{P}[\theta = 1 \mid x_i] = x_i$, and $\mathbb{P}[a_i = \theta \mid q_i = 1/2] = 1/2$ by (10),

$$\mathbb{P}[\theta = a_i \mid x_i] = \begin{cases} x_i & \text{if } x_i > 1/2 \\ 1 - x_i & \text{if } x_i < 1/2 \\ 1/2 & \text{if } x_i = 1/2. \end{cases}$$

Thus $\mathbb{P}[\theta = a_i \mid x_i] = \max\{x_i, 1 - x_i\}$. ∎

Let

$$x = \mathbb{P}[\theta = 1 \mid \bar{a}],$$

and note that x_i is a bounded martingale that converges a.s. to x. It thus follows from Claim 1 that conditioned on a_i taking both values infinitely often it holds that $x = 1/2$. Thus, to prove our theorem, we will show that the probability of $x = 1/2$ is zero. Accordingly, define the event

$$F^0 = \{x = 1/2\},$$

and for $\varepsilon > 0$ define the events

$$F_i^\varepsilon = \{x_i \in (1/2 - \varepsilon, 1/2 + \varepsilon)\}.$$

The event F_i^ε is the event that the public belief x_i is close to 1/2. Since the sequence $(x_i)_i$ converges a.s. to x, we have that

$$\lim_{i\to\infty} \mathbb{P}[F_0 \setminus F_\varepsilon^i] = 0 \tag{11}$$

for every $\varepsilon > 0$, and that

$$\lim_{\varepsilon\to 0} \limsup_{i\to\infty} \mathbb{P}[F_i^\varepsilon] = \mathbb{P}[F^0]. \tag{12}$$

Thus, to prove that $\mathbb{P}[F^0] = 0$—which, as we explained above, proves the claim—it suffices to show that the left hand side of the above expression vanishes.

To this end, let

$$b_i = b(s_i) \in \operatorname*{argmax}_{a\in A} \mathbb{P}[\theta = a \mid s_i]$$

be an optimal action chosen given agent i's private signal only. These are all chosen using the same function b, and so, since the private signals s_i are identically distributed (but not necessarily independently), the random variables b_i are also identically distributed.

Imagine that player i chooses b_i instead of a_i, whenever F_i^ε occurs. Then player i's gain in expected utility from this deviation is

$$\mathbb{P}[b_i = \theta, F_i^\varepsilon] - \mathbb{P}[a_i = \theta, F_i^\varepsilon].$$

We prove that the left-hand side of (12) vanishes by showing that

$$\lim_{\varepsilon\to 0} \limsup_{i\to\infty} \mathbb{P}[b_i = \theta, F_i^\varepsilon] - \mathbb{P}[a_i = \theta, F_i^\varepsilon] > 0,$$

and thus this is a profitable deviation for some ε small enough and i large enough, contradicting the assumption that $\bar{a}$ is an SLE.

To this end, we note that

$$\mathbb{P}[b_i = \theta, F_i^\varepsilon] \geq \mathbb{P}[b_i = \theta, F^0] - \mathbb{P}[b_i = \theta, F^0 \setminus F_i^\varepsilon],$$

since

$$F^0 \setminus (F^0 \setminus F_i^\varepsilon) = F^0 \cap F_i^\varepsilon \subseteq F_i^\varepsilon.$$

It thus follows by (11) that

$$\liminf_{i\to\infty} \mathbb{P}[b_i = \theta, F_i^\varepsilon] \geq \liminf_{i\to\infty} \mathbb{P}[b_i = \theta, F^0].$$

Now, conditioned on θ, the s_i's are i.i.d., and in particular mixing. Since F_0 is measurable in $\sigma(\bar{s})$, it follows that

$$\lim_{i\to\infty} \mathbb{P}[b_i = \theta, F^0 \mid \theta] = \mathbb{P}[b_i = \theta \mid \theta] \cdot \mathbb{P}[F^0 \mid \theta],$$

where the right-hand side does not depend on i, since the b_i's are identically distributed. Hence

$$\lim_{i\to\infty} \mathbb{P}[b_i = \theta, F^0] = \mathbb{P}[b_i = \theta] \cdot \mathbb{P}[F^0].$$

Since private signals are informative, it follows that $\mathbb{P}[b_i = \theta] > 1/2$, and so we have that

$$\liminf_{\varepsilon\to 0} \liminf_{i\to\infty} \mathbb{P}[b_i = \theta, F_i^\varepsilon] > \frac{1}{2}\mathbb{P}[F^0]. \tag{13}$$

Now,

$$\begin{aligned}\mathbb{P}[a_i = \theta \mid F_i^\varepsilon] &= \mathbb{E}\Big[\mathbb{P}[a_i = \theta \mid x_i] \mid F_i^\varepsilon\Big] \\ &= \mathbb{E}\Big[\max\{x_i, 1 - x_i\} \mid F_i^\varepsilon\Big],\end{aligned}$$

where the second equality is an application of Claim 2. Since $x_i \in (1/2 - \varepsilon, 1/2 + \varepsilon)$ conditioned on F_i^ε, we get that

$$\mathbb{P}[a_i = \theta, F_i^\varepsilon] < \left(\frac{1}{2} + \varepsilon\right) \cdot \mathbb{P}[F_i^\varepsilon].$$

Therefore, by (12),

$$\lim_{\varepsilon\to 0}\limsup_{i\to\infty} \mathbb{P}[a_i = \theta, F_i^\varepsilon] \leq \frac{1}{2}\cdot \mathbb{P}[F^0].$$

Therefore, in combination with (13), the expected profit from deviating from a_i to b_i on F_i^ε satisfies

$$\lim_{\varepsilon\to 0}\limsup_{i\to\infty} \mathbb{P}[b_i = \theta, F_i^\varepsilon] - \mathbb{P}[a_i = \theta, F_i^\varepsilon] > 0,$$

and thus this is a profitable deviation for some ε small enough and i large enough. Hence it follows that F^0 has probability zero, concluding the proof of Theorem 5.

C Mixing signals

Proof of Lemma 1. Let X be the indicator of the event B, and let Y_i be the indicator of the event $s_i \in S'$. Then we can write the expression in the lemma statement as

$$\mathbb{P}[s_i \in S', B \mid \theta] - \mathbb{P}[s_i \in S' \mid \theta]\cdot \mathbb{P}[B \mid \theta] = \mathrm{Cov}(X, Y_i).$$

Let $\eta_i \in \{-1, +1\}$ equal the sign of $\mathrm{Cov}(X, Y_i)$. Then

$$\left|\mathbb{P}[s_i \in S', B \mid \theta] - \mathbb{P}[s_i \in S' \mid \theta]\cdot \mathbb{P}[B \mid \theta]\right| = \mathrm{Cov}(X, \eta_i Y_i).$$

Thus, to prove our claim, we need to show that $\mathrm{Cov}(X, \eta_i Y_i) \leq \varepsilon$, except for finitely many i's.

Now, enumerate the agents $N = \{1, 2, 3, \ldots\}$ arbitrarily. Then

$$\sum_{i=1}^{n} \mathrm{Cov}(X, \eta_i Y_i) = \mathrm{Cov}\left(X, \sum_{i=1}^{n} \eta_i Y_i\right) \leq \sqrt{\mathrm{Var}(X)\cdot \mathrm{Var}\left(\sum_{i=1}^{n} Y_i\right)}. \tag{14}$$

To bound this expression, we first note that X is an indicator, and so $\mathrm{Var}(X) \leq 1$. Second,

$$\mathrm{Var}\left(\sum_{i=1}^{n} Y_i\right) = \sum_{i=1}^{n} \mathrm{Var}(Y_i) + \sum_{i=1}^{n}\sum_{j\neq i, j\leq n} \mathrm{Cov}(Y_i, Y_j).$$

The mixing property implies directly that for each i, $\mathrm{Cov}(Y_i, Y_j)$ is larger than ε for only finitely many j's. Hence

$$\sum_{j\neq i, j\leq n} \mathrm{Cov}(Y_i, Y_j) \leq f(n)$$

for some function f such that $\lim_{n\to\infty} f(n)/n = 0$. Applying all this back into (14) yields

$$\sum_{i=1}^{n} \mathrm{Cov}(X, \eta_i Y_i) \leq \sqrt{n + nf(n)}.$$

In particular, the average of $\mathrm{Cov}(X, \eta_i Y_i)$, as i ranges from 1 to n, satisfies

$$\frac{1}{n}\sum_{i=1}^{n} \mathrm{Cov}(X, \eta_i Y_i) \leq \sqrt{\frac{1}{n} + \frac{f(n)}{n}},$$

and in particular, since $\lim_{n\to\infty} f(n)/n = 0$, this average tends to zero. Thus, for each ε, each $\mathrm{Cov}(X, \eta_i Y_i) \leq \varepsilon$, except for finitely many values of i. ∎

D We must learn to agree

We say that an SLE satisfies *herding in probability* if there is a random variable a^* taking values in A such that

$$\lim_{i\to\infty} \mathbb{P}[a_i = a^*] = 1.$$

Here the limit is taken by arbitrarily identifying the agents with the set of natural numbers.

Theorem 4 is a consequence of the following, stronger statement that applies to herding in probability, rather than (almost sure) herding in which a cofinite set of agents chooses the same action.

Theorem 8. *In canonical* setting, every SLE that satisfies probability in herding also satisfies information diffusion.*

Let

$$p = \mathbb{P}[a^* = \theta].$$

It follows from herding in probability that $\lim_i \mathbb{P}[a_i = a^*] = 1$, and so

$$\begin{aligned}\lim_i \mathbb{P}[a_i = \theta] &= \lim_i \mathbb{P}[a_i = \theta, a_i = a^*] + \mathbb{P}[a_i = \theta, a_i \neq a^*] \\ &= \lim_i \mathbb{P}[a^* = \theta, a_i = a^*] + \mathbb{P}[a_i = \theta, a_i \neq a^*] \\ &= p.\end{aligned}$$

Assume by contradiction that

$$p \le 1 - \beta - 2\varepsilon$$

for some $\varepsilon > 0$.

As in the proof of Theorem 5, let

$$b_i = b(s_i) \in \operatorname*{argmax}_{a \in A} \mathbb{P}[\theta = a \mid s_i]$$

be an optimal action chosen given agent i's private signal only. These are all chosen using the same function b, and so, since the private signals s_i are identically distributed (but not necessarily independently), the random variables b_i are also identically distributed. Let B_i be the event that $\mathbb{P}[b_i = \theta \mid s_i] > 1 - \beta - \varepsilon$. Since the b_i's are identically distributed, all of the events B_i have the same probability. Furthermore, this probability is positive, by our assumption on the support of the private signals.

Imagine that agent i deviates and chooses b_i whenever B_i occurs, and otherwise follows a_i. Then her expected gain in utility is

$$\mathbb{P}[b_i = \theta, B_i] - \mathbb{P}[a_i = \theta, B_i].$$

To bound the first term, we note that, by the definition of B_i,

$$\mathbb{P}[b_i = \theta, B_i] \ge (1 - \beta - \varepsilon)\mathbb{P}[B_i].$$

To bound the second term, we write

$$\begin{aligned}\mathbb{P}[a_i = \theta, B_i] &= \mathbb{P}[a_i = \theta, a_i = a^*, B_i] + \mathbb{P}[a_i = \theta, a_i \ne a^*, B_i] \\ &= \mathbb{P}[a^* = \theta, a_i = a^*, B_i] + \mathbb{P}[a_i = \theta, a_i \ne a^*, B_i]\end{aligned}$$

Since $\bar{a}$ satisfies herding in probability, $\lim_i \mathbb{P}[a_i = a^*] = 1$, and so it follows that

$$\limsup_i \mathbb{P}[a_i = \theta, B_i] = \limsup_i \mathbb{P}[a^* = \theta, B_i].$$

Since private signals are conditionally mixing, it follows from Lemma 1 that

$$\limsup_i \mathbb{P}[a^* = \theta, B_i] = \mathbb{P}[a^* = \theta]\mathbb{P}[B_i] = p \cdot \mathbb{P}[B_i],$$

where the right-hand side does not depend on i, since the events B_i all have the same probability. We have thus shown that

$$\limsup_i \mathbb{P}[a_i = \theta, B_i] = p \cdot \mathbb{P}[B_i],$$

Combining the bounds on the two terms we get that the expected gain in utility is

$$\liminf_i \mathbb{P}[b_i = \theta, B_i] - \mathbb{P}[a_i = \theta, B_i] \geq (1 - \beta - \varepsilon - p)\mathbb{P}[B_i].$$

Since we assumed that $p \leq 1 - \beta - 2\varepsilon$ we have that this is at least $\varepsilon\mathbb{P}[B_i]$, and in particular positive. Thus $\bar{a}$ is not an equilibrium, as for some i large enough player i would have a profitable deviation. This completes the proof of Theorem 8.

E Proof of Theorem 6

This proof is essentially a recasting of the proof of Proposition 2.1 in Rosenberg et al. (2009) to our language and notation.

Fix an agent i. The case that $\delta = 0$ or T_i is finite is immediate. We thus assume henceforth that $\delta > 0$ and $|T_i| = \infty$.

Let

$$v_i = \max_{a \in A} \mathbb{E}[u(a, \theta) \mid \Bbbk_i, s_i]$$

be the maximum expected utility agent i can guarantee given what she (asymptotically) knows at the end of the game.

Fix $(\Bbbk_i, s_i)$ and $\varepsilon > 0$, and let $U, W \subseteq A$ be the sets of actions given by

$$U = \left\{a \in A : \mathbb{E}[u(a, \theta) \mid \Bbbk_i, s_i] > v_i - \varepsilon\right\}$$

and

$$W = \left\{b \in A : \mathbb{E}[u(b, \theta) \mid \Bbbk_i, s_i] < v_i - 3\varepsilon\right\}.$$

That is, U is the set of actions that is ε-optimal, and W is the set of actions that is 3ε-suboptimal—conditioned on the information available to the player at the end of the game.

Note that the sets U and W are open and disjoint and that utilities are continuous. It therefore follows from the martingale convergence theorem that, for any $\eta > 0$ and $t \in T_i$

large enough, it holds for every $a \in U$ and $b \in W$ that

$$\mathbb{P}\left[\mathbb{P}[u(a,\theta) > u(b,\theta) + \varepsilon \mid \Bbbk_{i,t}, s_i] > 1 - \eta \mid \Bbbk_i, s_i\right] = 1.$$

That is, for large enough t, the agent will almost surely assign high probability to the event that any action $a \in U$ yields at least ε more utility than any $b \in W$.

It follows that choosing any $b \in W$ will, for large enough t, result in an expected utility loss of at least $\varepsilon \cdot (1-\eta) \cdot (1-\delta)$ in the subgame starting at t, which, for η small enough, is greater than $\delta \cdot \max_{a,\theta} u(a,\theta)$, and thus greater than the continuation utility of any strategy. It follows that, in equilibrium and conditioned on almost every $(\Bbbk_i, s_i)$, the agent eventually stops choosing actions in W. Since this holds for every ε, it follows that every limit point of actions taken by the agents almost surely maximizes her expected utility, conditioned on $(\Bbbk_i, s_i)$.